arXiv:1407.5011v2 [cs.DS] 26 Apr 2015

A simple and optimal ancestry labeling scheme for trees

Søren Dahlgaard*, Mathias Bæk Tejs Knudsen*†, and Noy Rotbart

Department of Computer Science, University of Copenhagen
Universitetsparken 5, 2100 Copenhagen
{soerend,knudsen,noyro}@di.ku.dk

April 28, 2015

Abstract

We present a $\lg n + 2\lg\lg n + 3$ ancestry labeling scheme for trees. The problem was first presented by Kannan et al. [STOC 88'] along with a simple $2\lg n$ solution. Motivated by applications to XML files, the label size was improved incrementally over the course of more than 20 years by a series of papers. The last, due to Fraigniaud and Korman [STOC 10'], presented an asymptotically optimal $\lg n + 4\lg\lg n + O(1)$ labeling scheme using non-trivial tree-decomposition techniques. By providing a framework generalizing interval based labeling schemes, we obtain a simple, yet asymptotically optimal solution to the problem. Furthermore, our labeling scheme is attained by a small modification of the original $2\lg n$ solution.

1 Introduction

The concept of *labeling schemes*, introduced by Kannan, Naor and Rudich [16], is a method to assign bit strings, or labels, to the vertices of a graph such that a query between vertices can be inferred directly from the assigned labels, without using a centralized data structure. A labeling scheme for a family of graphs $\mathcal{F}$ consists of an *encoder* and a *decoder*. Given a graph $G \in \mathcal{F}$, the encoder assigns labels to each node in G, and the decoder can infer the query given only a set of labels. The main quality measure for a labeling scheme is the size of the largest label size it assigns to a node of any graph of the entire family. One of the most well studied questions in the context of labeling schemes is the *ancestry problem*. An ancestry labeling scheme for the family of rooted trees of n nodes $\mathcal{F}$ assigns labels to a tree $T \in \mathcal{F}$ such that given the labels $\ell(u), \ell(v)$ of any two nodes $u, v \in T$, one can determine whether u is an *ancestor* of v in T.

Improving the label size for this question is highly motivated by XML search engines. An XML document can be viewed as a tree and queries over such documents amount to testing ancestry relations between nodes of these trees [2, 11, 10]. Search engines process queries using an index structure summarizing the ancestor relations. It is imperative to the performance of such engines,

*Research partly supported by Thorup's Advanced Grant from the Danish Council for Independent Research under the Sapere Aude research career programme.

†Research partly supported by the FNU project AlgoDisc - Discrete Mathematics, Algorithms, and Data Structures.

that as much as possible of this index can reside in main memory. The big size of web data thus implies that even a small reduction in label size may significantly improve the memory cost and performance. A more detailed explanation can be found in [1].

One solution to this problem, dating back to at least '74 by Tarjan [19] and used in the seminal paper in '92 by Kannan et al. [16] is the following: Given an n node tree T rooted in r, perform a DFS traversal and for each $u \in T$ let $\mathrm{dfs}(u)$ be the index of u in the traversal. Assign the label of u as the pair $\ell(u) = (\mathrm{dfs}(u), \mathrm{dfs}(v))$, where v is the descendant of u of largest index in the DFS traversal. Given two labels $\ell(u) = (\mathrm{dfs}(u), \mathrm{dfs}(v))$ and $\ell(w)$, u is an ancestor of w iff $\mathrm{dfs}(u) \leqslant \mathrm{dfs}(w) \leqslant \mathrm{dfs}(v)$. In other words the label $\ell(u)$ represents an interval, and ancestry is determined by interval containment. Since every DFS index is a number in $\{1, \ldots, n\}$ The size of a label assigned by this labeling scheme is at most $2 \lg n$.[1]

This labeling scheme, denoted Classic from hereon, was the first in a long line of research to minimize the label size to $1.5 \lg n$ [3], $\lg n + O(\lg n / \lg \lg n)$ [20][2], $\lg n + O(\sqrt{\lg n})$ [8, 1] and finally $\lg n + 4 \lg \lg n + O(1)$ [14], essentially matching a lower bound of $\lg n + \lg \lg n - O(1)$ [4]. Additional results for this labeling scheme are a $\lg n + O(\lg \delta)$ labeling scheme for trees of depth at most δ [13] and investigation on a dynamic variant with some pre-knowledge on the structure of the constructed tree [10]. The asymptotically optimal labeling scheme [14] also implies the existence of a *universal poset* of size $O(n^k \lg^{4k} n)$.

Labeling schemes for other functions were considered. Among which are adjacency [16, 9, 7], routing [12, 20], nearest common ancestor [5, 6] connectivity [18], distance [15], and flow [17].

1.1 Our contribution

We present a *simple* ancestry labeling scheme of size $\lg n + 2 \lg \lg n + 3$. Similarly to [1, 14] our labeling scheme is based on assigning an interval to each node of the tree. Our labeling scheme can be seen as an extension of the Classic scheme described above, with the key difference that rather than storing the exact size of the interval, we store only an approximation thereof. In order to store only this approximation, our scheme assigns intervals larger than needed, forcing us to use a range larger than $1, \ldots, n$. Our main technical contribution is to minimize the label size by balancing the approximation ratio with the size of the range required to accommodate the resulting labels. While it is a challenge to prove the mathematical properties of our labeling scheme, describing it can be done in a few lines of pseudocode.

The simplicity of our labeling scheme contrast the labeling scheme of [14] , which relies among other things, on a highly nontrivial tree decomposition that must precede it's encoding. As a concrete example, our encoder can be implemented using a single DFS traversal, and results in a label size of $\lg n + 2 \lg \lg n + O(1)$ bits compared to $\lg n + 4 \lg \lg n + O(1)$. However, for trees of constant depth, our scheme has size $\lg n + \lg \lg n + O(1)$, which is worse than the bound achieved in [13].

The paper is structured as follows: In Section 2 we present a general framework for describing ancestry labeling schemes based on the notion of *left-including intervals*, which we introduce in Definition 1. Lemmas 1 and 2 show the correctness of all labeling schemes which construct labels by assigning *left-including intervals* to nodes. We illustrate the usefulness of this framework by using it to describe Classic in Section 3. Finally, In Section 4 we use the framework to construct our

[1] Throughout this paper we use $\lg n$ to denote the base 2 logarithm $\lg_2 n$.

[2] The paper discussed proves this bound on routing, which can be used to determine ancestry.

new approximation-based ancestry labeling scheme.

1.2 Preliminaries

We use the notation $[n] = \{0, 1, \ldots, n-1\}$ and denote the concatenation of two bit strings a and b by $a \circ b$. Let $T = (V, E)$ be a tree rooted in r, where $u, v \in V$. The node u is an *ancestor* of v if u lies on the unique path from the root to v, and we call v a *descendant* of u iff u is an ancestor of v. We denote the subtree rooted in u as T_u, i.e. the tree consisting of all descendants of u, and stress that a node is both an ancestor and descendant of itself. The encoding and decoding running time are described under the word-RAM model.

We denote the *interval* assigned to a node u by $I(u) = [a(u), b(u)]$, where $a(u)$ and $b(u)$ denote the lower and upper part of the interval, respectively. We also define $\overline{a}(u)$ and $\overline{b}(u)$ to be the maximum value of $a(v)$ respectively $b(v)$, where v is a descendant of u (note that this includes u itself). We will use the following notion:

Definition 1. Let T be a rooted tree and I an interval assignment defined on $V(T)$. We say that the interval assignment I is *left-including* if for each $u, v \in T$ it holds that u is an ancestor of v iff $a(v) \in I(u)$.

In contrast to Definition 1, the literature surveyed [16, 3, 8, 14] considers intervals where u is an ancestor of v iff $I(v) \subseteq I(u)$, i.e. the interval of a descendant node is fully contained in the interval of the ancestor. This distinction is amongst the unused leverage points which we will use to arrive at our new labeling scheme.

2 A framework for interval based labeling schemes

In this section we introduce a framework for assigning intervals to tree nodes. We will see in Sections 3 and 4 how this framework can be used to describe ancestry labeling schemes. The framework relies heavily on the values defined in Section 1.2, namely $a(u), b(u), \overline{a}(u), \overline{b}(u)$. An illustration of these values is found in Figure 1 below. The interval $[\overline{a}(u), \overline{b}(u)]$ can be seen as a slack interval from which $b(u)$ can be chosen. This will prove useful in Section 4.

The following lemmas contain necessary and sufficient conditions for interval assignments satisfying the left inclusion property.

Lemma 1. *Let T be a rooted tree and I a left-including interval assignment defined on $V(T)$. Then the following is true:*

1. *For each $u \in T$, $b(u) \geqslant \overline{a}(u)$.*
2. *For each $u \in T$ and $v \in T_u \backslash \{u\}$ a descendant of u, $a(v) > a(u)$.*
3. *For each $u \in T$, $[a(u), \overline{b}(u)] = \bigcup_{v \in T_u} I(v) = \bigcup_{v \in T_u} [a(v), b(v)]$*
4. *For any two distinct nodes $u, v \in T$ such that u is not an ancestor of v and v is not an ancestor of u the intervals $[a(u), \overline{b}(u)]$ and $[a(v), \overline{b}(v)]$ are disjoint.*

Proof. We prove that if the set of intervals is left-including each of the four statements must hold. We finish the proof of each statement with a $\diamond$.

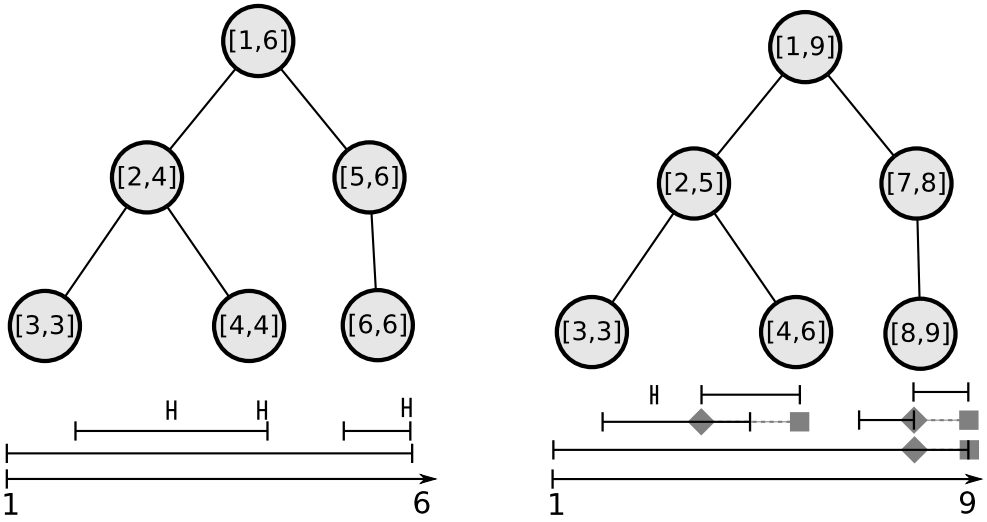

Figure 1: Two examples of left-including interval assignments to a tree. Left: a left-including assignment as used for CLASSIC in the introduction corresponding to $b(u) = \overline{a}(u)$. Right: a different left-including assignment for the same tree. For internal nodes where $\overline{a}(u)$ and $\overline{b}(u)$ do not coincide with $b(u)$, we have marked these by a gray diamond and square respectively.

1. Let $u \in T$. For any descendant $v \in T_u$ of u we know that $a(v) \in [a(u), b(u)]$. In particular $a(v) \leqslant b(u)$. Since this holds for any $v \in T_u$ we conclude that $\overline{a}(v) = \max_{v \in T_u} \{a(v)\} \leqslant b(u)$. $\diamond$

2. Let $u \in T$ and $v \in T_u \backslash \{u\}$ be a descendant of u. Since v is a descendant of u we see that $a(v) \in [a(u), b(u)]$ and in particular $a(v) \geqslant a(u)$. Since u is a not a descendant of v we know that $a(u) \notin [a(v), b(v)]$ and in particular $a(u) \neq a(v)$. Hence $a(v) > a(u)$. $\diamond$

3. Let $u \in T$. For any $v \in T_u$ $a(v) \in I(u)$ and hence $I(u)$ and $I(v)$ must have a non-empty intersection. Hence $\bigcup_{v \in T_u} I(v)$ must be the union of overlapping intervals, and therefore it is an interval. Since $a(v) \geqslant a(u)$ for any $v \in T_u$ by part 2, the minimum value in the interval must be $a(u)$. The maximal value must be $\max_{v \in T_u} \{b(v)\}$ which is $\overline{b}(u)$ by definition. Hence the interval must be $[a(u), \overline{b}(u)]$. $\diamond$

4. Let $u, v \in T$ be given as in the statement of the condition. Assume for the sake of contradiction that $[a(u), \overline{b}(u)]$ and $[a(v), \overline{b}(v)]$ are not disjoint. By part 3 this is equivalent to stating that the following sets are not disjoint:

$$\bigcup_{w \in T_u} [a(w), b(w)], \bigcup_{z \in T_v} [a(z), b(z)]$$

Hence there must exist $w \in T_u, z \in T_v$ such that $[a(w), b(w)]$ and $[a(z), b(z)]$ are overlapping. Assume wlog that $a(w) \leqslant a(z)$. Since the intervals overlap $b(w) \geqslant a(z)$ and hence $a(z) \in [a(w), b(w)]$ and since the set of intervals is left-including this implies that w is an ancestor of z. But this cannot be the case if none of u, v is an ancestor of the other. Contradiction. Hence the assumption was wrong and the intervals are indeed disjoint. $\diamond$ $\square$

Lemma 2. *Let T be a rooted tree and I an interval assignment defined on $V(T)$. If the following conditions are satisfied, then I is a left-including interval assignment.*

i For each $u \in T$, $b(u) \geqslant \overline{a}(u)$.

ii For each $u \in T$ and $v \in T$ a child of u, $a(v) > a(u)$.

iii For any two siblings $u, v \in T$ the intervals $[a(u), \overline{b}(u)]$ and $[a(v), \overline{b}(v)]$ are disjoint.

Proof. We prove that if the three conditions hold the set of intervals are left-including. First we prove that condition *ii* implies condition 2 from Lemma 1:

For each $u \in T$ and $v \in T_u \backslash \{u\}$ a descendant of u, $a(v) > a(u)$.
Let $u \in T, v \in T_u \backslash \{u\}$ and consider the path from u to v, say $u = v_0 \to v_1 \to \ldots \to v_k = v$ for positive integer k. Then v_i is the child of v_{i-1} for each $i = 1, 2, \ldots, k$. Therefore, by condition *ii*:

$$a(u) = a(v_0) < a(v_1) < \ldots < a(v_k) = a(v)$$

Which proves condition 2.

Fix distinct $u, v \in T$. We will show that u is an ancestor of v if and only if $a(v) \in I(u)$. We do this by considering three cases where u is an ancestor of v, v is an ancestor of u and $v \neq u$, and none is an ancestor of the other respectively.

Case 1: u is an ancestor of v. By condition *i* $b(u) \geqslant \overline{a}(u)$. Since $v \in T_u$ this implies $b(u) \geqslant a(v)$. By condition 2 $a(v) > a(u)$. Therefore $a(u) < a(v) \leqslant b(u)$ and $a(v) \in I(u)$.

Case 2: v is an ancestor of u and $v \neq u$. By condition 2 $a(u) > a(v)$. Hence $a(v)$ is smaller than any value in the interval $[a(u), b(u)]$, and therefore $a(v) \notin I(u)$.

Case 3: None of u, v is an ancestor of the other. By assumption u, v must have a nearest common ancestor $w \in T$ which is neither u or v. Let u' be the child of w that lies on the path from w to u. Note that if u is the child of w then $u = u'$. Define v' analogously. By condition *iii* $[a(u'), \overline{b}(u')]$ and $[a(v'), \overline{b}(v')]$ are disjoint. Hence either $\overline{b}(u') < a(v')$ or $\overline{b}(v') < a(u')$. Since both cases are handled similarly we will show how to handle the case $\overline{b}(u') < a(v')$. By condition 2 $a(v) > a(v')$ and hence $a(v) > \overline{b}(u')$. Recall that $\overline{b}(u')$ is the maximum of $b(z)$ where $z \in T_{u'}$. Since $T_u \subseteq T_{u'}$ this implies that $\overline{b}(u) \leqslant \overline{b}(u')$. Hence $a(v) > \overline{b}(u)$ as well and the intervals $[a(v), \overline{b}(v)]$ and $[a(u), \overline{b}(u)]$ are disjoint. Hece $I(u)$ and $I(v)$ must be disjoint as well and $a(v) \notin I(u)$. □

2.1 The framework

We now consider a general approach for creating left-including interval assignments. For a node $u \in T$ and a positive integer t we define the procedure ASSIGN(u, t) that assigns intervals to T_u recursively and in particular, assigns $a(u) = t$. For pseudocode of the procedure see Algorithm 1.

Algorithm 1 Assigning intervals to all nodes in the subtree T_u rooted at u ensuring $a(u) = t$.

```
1: procedure ASSIGN(u, t)
2:     (a(u), ā(u), b(u), b̄(u)) ← (t, t, t, t)
3:     for v ∈ children(u) do
4:         ASSIGN(v, b̄(u) + 1)
5:         (ā(u), b̄(u)) ← (ā(v), b̄(v))
6:     Assign b(u) such that b(u) ⩾ ā(u).
7:     b̄(u) ← max{b(u), b̄(u)}
```

Algorithm 1 provides a general framework for assigning intervals using a depth-first traversal. We can use it to design an actual interval assignment by specifying: (1) the way we choose $b(u)$, and (2) the order in which the children are traversed. These specifications correspond to line 6

and line 3, respectively, and determine entirely the way the intervals are assigned. It may seem counter-intuitive to pick $b(u) > \bar{a}(u)$, but we will show that doing so in a systematic way, we are able to describe the interval using fewer bits by limiting the choices for $b(u)$. In the remainder of this paper, we will see how these two decisions impact also the label size, and produce our claimed labeling scheme.

We now show that any ordering of the children and any way of choosing $b(u)$ satisfying $b(u) \geqslant \overline{a}(u)$ generates a left-including interval assignment.

Lemma 3. *Let T be a tree rooted in r. After running Algorithm 1 with* ASSIGN$(r, 0)$ *the values of $\overline{a}(u), \overline{b}(u)$ are correct, i.e. for all $u \in T$:*

$$\overline{a}(u) = \max_{v \in T_u} \{a(v)\}, \quad \overline{b}(u) = \max_{v \in T_u} \{b(v)\}.$$

Proof. We prove the claim by induction on the size of the subtree $|T_u|$, and prove that the invariant holds when the function call ASSIGN(u, t) is terminated.This is sufficient, since none of the involved variables are altered afterwards.

As base case $|T_u| = 1$ and u is a leaf and the claim holds trivially. Assume next that the claim holds for $|T_u| \leqslant m$ for some integer $m \geqslant 1$, and let $u \in T$ be of size $|T_u| = m + 1$, and $v_1, \ldots, v_k$ be the children of u in the order they are processed on Line 3. First, note that the size of each $|T_{v_1}| \ldots |T_{v_k}| \leqslant m$. It follows that $\overline{a}(v_i)$ and $\overline{b}(v_i)$ are correct for the nodes v_i where $1 \leqslant i \leqslant v_k$. Line 5 guarantees that $\overline{b}(u)$ is at least as large as the largest $\overline{b}(v_i)$, and that $\overline{a}(u)$ is at least as large as the largest $\overline{a}(v_i)$. □

The following Lemma is useful for showing several properties in the framework.

Lemma 4. *Let u be a node in a tree T with children $v_1 \ldots v_k$. After running Algorithm 1 with parameters* ASSIGN$(r, 0)$ *where $v_1 \ldots v_k$ are processed in that order, the following properties hold:*

1. $\overline{b}(u) - a(u) + 1 = \left(\sum_{i=1}^{k} \overline{b}(v_i) - a(v_i) + 1\right) + 1.$
2. $\overline{a}(u) - a(u) + 1 = \overline{a}(v_k) - a(v_k) + \left(\sum_{i=1}^{k-1} \overline{b}(v_i) - a(v_i) + 1\right) + 1.$

Proof. By the definition of ASSIGN we see that for all $i = 1, \ldots, k-1$, $a(v_{i+1}) = 1 + \overline{b}(v_i)$. Furthermore $a(v_1) = a(u) + 1$ and $\overline{b}(v_k) = \overline{b}(u)$. Hence:

$$\begin{aligned}
\overline{b}(u) - a(u) + 1 &= \overline{b}(v_k) - a(v_1) + 2 \\
&= \left(\sum_{i=2}^{k} \overline{b}(v_i) - \overline{b}(v_{i-1})\right) + \overline{b}(v_1) - a(v_1) + 2 \\
&= \left(\sum_{i=1}^{k} \overline{b}(v_i) - a(v_i) + 1\right) + 1.
\end{aligned}$$

The second equality follows by the same line of argument. □ □

Theorem 1. *Let T be a tree rooted in r. After running Algorithm 1 with parameters* ASSIGN$(r, 0)$ *the set of intervals produced are left-including.*

Proof. Consider any node $u \in T$ and a call ASSIGN(u, t). We will prove each of the conditions of Lemma 2, which implies the theorem.

- i This condition is trivially satisfied by line 6.
- iii First, observe that any interval assigned to a node w by a call to ASSIGN(v, t) has $a(w) \geqslant t$, and by i it has $\overline{b}(w) \geqslant b(w) \geqslant a(w)$. Let $v_1, \ldots, v_k$ be the children in the order of the for loop in line 3. By lines 2, 4 and 5 we have $a(v_1) = a(u) + 1, a(v_2) = \overline{b}(v_1) + 1, a(v_3) = \overline{b}(v_2) + 1, \ldots, a(v_k) = \overline{b}(v_{k-1}) + 1$, thus the condition is satisfied.
- ii The first child v of u has $a(v) = t + 1 = a(u) + 1$. By the same line of argument as in *iii* we see that all other children w of u must have $a(w) > a(v) = a(u) + 1$. □

□

3 The classic ancestry labeling scheme

To get acquainted with the framework of Section 2, we use it to redefine CLASSIC, the labeling scheme introduced in Section 1.

Let T be a tree rooted in r. We first modify the function ASSIGN to create ASSIGN-CLASSIC such that the intervals $I(u) = [a(u), b(u)]$ correspond to the intervals of the CLASSIC algorithm described in the introduction. To do this we set $b(u) = \overline{a}(u)$ in line 6 and traverse the children in any order in line 3. We note that there is a clear distinction between an algorithm such as ASSIGN-CLASSIC and an encoder. This distinction will be more clear in Section 4. We will need the following lemma to describe the encoder.

Lemma 5. *After* ASSIGN-CLASSIC(u, t) *is called the following invariant is true:*

$$\overline{b}(u) - a(u) + 1 = |T_u|$$

Proof. We prove the claim by induction on $|T_u|$. When $|T_u| = 1$ u is a leaf and hence $\overline{b}(u) = a(u) = t$ and the claim holds.

Let $|T_u| = m > 1$ and assume that the claim holds for all nodes with subtree size $< m$. Let $v_1, \ldots, v_k$ be the children of u. By Lemma 4 and the induction hypothesis we have:

$$\begin{aligned} \overline{b}(u) - a(u) + 1 &= \left(\sum_{i=1}^{k} \overline{b}(v_i) - a(v_i) + 1 \right) + 1 \\ &= \left(\sum_{i=1}^{k} |T_{v_i}| \right) + 1 = |T_u| \,. \end{aligned}$$

This completes the induction. □ □

Description of the encoder: Let T be an n-node tree rooted in r. We first invoke a call to ASSIGN-CLASSIC$(r, 0)$. By Lemma 5 we have $\overline{b}(r) - a(r) + 1 = n$ and this implies $0 \leqslant a(u), b(u) \leqslant n - 1$ for every $u \in T$. Let x_u and y_u be the encoding of $a(u)$ and $b(u)$ using exactly[3] $\lceil \lg n \rceil$ bits respectively. We set the label of u to be the concatenation of the two bitstrings, i.e. $\ell(u) = x_u \circ y_u$.

[3]This can be accomplished by padding with zeros if necessary.

Description of the decoder: Let $\ell(u)$ and $\ell(v)$ be the labels of the nodes u and v in a tree T. By the definition of the encoder, the labels have the same size and it is $2z$ for some integer $z \geqslant 1$. Let $\ell(u) = x_u \circ y_u$ where x_u and y_u are the first and last z bits of $\ell(u)$ respectively. Let a_u and b_u the integers from $[2^z]$ corresponding to the bit strings x_u and y_u respectively. We define a_v and b_v analogously. The decoder responds `True`, i.e. that u is the ancestor of v, iff $a_v \in [a_u, b_u]$.

The correctness of the labeling scheme follows from Theorem 1 and the description of the decoder.

4 An approximation-based approach

In this section we present the main result of this paper:

Theorem 2. *There exist an ancestry labeling scheme of size* $\lceil \lg n \rceil + 2 \lceil \lg \lg n \rceil + 3$.

To prove this theorem, we use the framework introduced in Section 2. The barrier in reducing the size of the Classic labeling scheme is that the number of different intervals that can be assigned to a node is $\Theta(n^2)$. It is impossible to encode so many different intervals without using at least $2 \lg n - O(1)$ bits. The challenge is therefore to find a smaller set of intervals $I(u) = [a(u), b(u)]$ to assign to the nodes. First, note that Lemma 1 points 2 and 4 imply that any two nodes u, v must have $a(u) \neq a(v)$. By considering the n node tree T rooted in r where r has $n-1$ children, we also see that there must be at least $n-1$ different values of $b(u)$ (by Lemma 1 point 4). One might think that this implies the need for $\Omega(n^2)$ different intervals. This is, however, not the case. We consider a family of intervals, such that $a(u) = O(n)$ and the size of each interval, $b(u) - a(u) + 1$, comes from a much smaller set, S. Since there are $O(n\,|S|)$ such intervals we are able to encode them using $\lg n + \lg |S| + O(1)$ bits.

We now present a modification of Assign called Assign-New. Calling Assign-New$(r, 0)$ on an n-node tree T with root r will result in each $a(u) \in [2n]$ and $b(u) \in S$, where S is given by:

$$S = \left\{ \left\lfloor (1+\varepsilon)^k \right\rfloor \mid k \in \left[4 \lceil \lg n \rceil^2 \right] \right\} , \tag{1}$$

where ε is the unique solution to the equation $\lg(1+\varepsilon) = (\lceil \lg n \rceil)^{-1}$. First, we examine some properties of S:

Lemma 6. *Let S be defined as in (1). For every $m \in \{1, 2, \ldots, 2n\}$ there exists $s \in S$ such that:*

$$m \leqslant s < m(1+\varepsilon) .$$

Furthermore, $s = \left\lfloor (1+\varepsilon)^k \right\rfloor$ for some $k \in \left[4 \lceil \lg n \rceil^2 \right]$, and both s and k can be computed in $O(1)$ time.

Proof. Fix $m \in \{1, 2, \ldots, 2n\}$. Let k be the largest integer such that $(1+\varepsilon)^{k-1} < m$. Equivalently, k is the largest integer such that:

$$k - 1 < \frac{\lg m}{\lg(1+\varepsilon)} = (\lg m) \cdot \lceil \lg n \rceil .$$

In other words we choose k as $\lceil (\lg m) \cdot \lceil \lg n \rceil \rceil$ and note that k is computed in $O(1)$ time. Since $\lg m \leqslant \lg(2n) \leqslant 2 \lg n$:

$$k \leqslant \lceil 2(\lg n) \cdot \lceil \lg n \rceil \rceil \leqslant 2 \lceil \lg n \rceil^2 < 4 \lceil \lg n \rceil^2 .$$

By setting $s = \left\lfloor (1+\varepsilon)^k \right\rfloor$ we have $s \in S$. By the definition of k we see that $(1+\varepsilon)^k \geqslant m$ and thus also $s \geqslant m$. Similarly:

$$m(1+\varepsilon) > (1+\varepsilon)^{k-1} \cdot (1+\varepsilon) = (1+\varepsilon)^k \geqslant s.$$

This proves that $s \in S$ satisfies the desired requirement. Furthermore s can be computed in $O(1)$ time by noting that:

$$s = \left\lfloor (1+\varepsilon)^k \right\rfloor = \left\lfloor 2^{\lg(1+\varepsilon)k} \right\rfloor = \left\lfloor 2^{\lceil \lg n \rceil^{-1} \cdot k} \right\rfloor.$$

□ □

We now define Assign-New by modifying Assign in the following two ways. First, we specify the order in which the children are traversed in line 3. This is done in non-decreasing order of their subtree size, i.e. we iterate $v_1, \ldots, v_k$, where $|T_{v_1}| \leqslant \ldots \leqslant |T_{v_k}|$. Second, we choose $b(u)$ in Line 6 as the smallest value, such that $b(u) \geqslant \overline{a}(u)$ and $b(u) - a(u) + 1 \in S$. This is done by using Lemma 6 with $m = \overline{a}(u) - a(u) + 1$ and setting $b(u) = a(u) + s - 1$. In order to do this we must have $m \leqslant 2n$. To do this, we show the following lemma corresponding to Lemma 5 in Section 3.

Lemma 7. *After* Assign-New(u, t) *is called the following invariants holds:*

$$\overline{a}(u) - a(u) + 1 \leqslant |T_u|\,(1+\varepsilon)^{\lfloor \lg |T_u| \rfloor} \tag{2}$$

$$\overline{b}(u) - a(u) + 1 \leqslant |T_u|\,(1+\varepsilon)^{\lfloor \lg |T_u| \rfloor + 1} \tag{3}$$

Proof. We prove the claim by induction on $|T_u|$. When $|T_u| = 1$, u is a leaf, so $\overline{b}(u) = \overline{a}(u) = a(u) = t$ and the claim holds.

Now let $|T_u| = m > 1$ and assume that the claim holds for all nodes with subtree size $< m$. Let $v_1, \ldots, v_k$ be the children of u such that $|T_{v_1}| \leqslant \ldots \leqslant |T_{v_k}|$. First, we show that (2) holds. By Lemma 4 we have the following expression for $\overline{a}(u) - a(u) + 1$:

$$\overline{a}(u) - a(u) + 1 = (\overline{a}(v_k) - a(v_k) + 1) + \left(\sum_{i=1}^{k-1} \overline{b}(v_i) - a(v_i) + 1 \right) + 1. \tag{4}$$

It follows from the induction hypothesis that:

$$\overline{a}(v_k) - a(v_k) + 1 \leqslant |T_{v_k}|\,(1+\varepsilon)^{\lfloor \lg |T_{v_k}| \rfloor} \leqslant |T_{v_k}|\,(1+\varepsilon)^{\lfloor \lg |T_u| \rfloor}. \tag{5}$$

Furthermore, by the ordering of the children, we have $\lg |T_{v_i}| \leqslant \lg |T_u| - 1$ for every $i = 1, \ldots, k-1$. Hence:

$$\overline{b}(v_i) - a(v_i) + 1 \leqslant |T_{v_i}|\,(1+\varepsilon)^{\lfloor \lg |T_{v_i}| \rfloor + 1} \leqslant |T_{v_k}|\,(1+\varepsilon)^{\lfloor \lg |T_u| \rfloor}. \tag{6}$$

Inserting (5) and (6) into (4) proves invariant (2).

Since $\overline{b}(u) = \max\left\{\overline{b}(v_k), b(u)\right\}$ we only need to upper bound $\overline{b}(v_k) - a(u) + 1$ and $b(u) - a(u) + 1$. First we note that since $b(u)$ is chosen smallest possible such that $b(u) \geqslant \overline{a}(u)$ and $b(u) - a(u) + 1 \in S$, it is guaranteed by Lemma 6 that:

$$b(u) - a(u) + 1 < (1+\varepsilon)\,(\overline{a}(u) - a(u) + 1) \leqslant |T_u|\,(1+\varepsilon)^{\lfloor \lg |T_u| \rfloor + 1}$$

Hence we just need to upper bound $\overline{b}(v_k) - a(u) + 1$. First we note that just as in (4):

$$\overline{b}(v_k) - a(u) + 1 = \left(\sum_{i=1}^{k} \overline{b}(v_i) - a(v_i) + 1 \right) + 1 \tag{7}$$

By the induction hypothesis, for every $i = 1, \ldots, k$:

$$\overline{b}(v_i) - a(v_i) + 1 \leqslant |T_{v_i}| \, (1+\varepsilon)^{\lfloor \lg |T_{v_i}| \rfloor + 1} \leqslant |T_{v_i}| \, (1+\varepsilon)^{\lfloor \lg |T_u| \rfloor + 1} \tag{8}$$

Inserting (8) into (7) gives the desired:

$$\overline{b}(v_k) - a(u) + 1 \leqslant 1 + \sum_{i=1}^{k} |T_{v_i}| \, (1+\varepsilon)^{\lfloor \lg |T_u| \rfloor + 1} \leqslant |T_u| \, (1+\varepsilon)^{\lfloor \lg |T_u| \rfloor + 1}.$$

This completes the induction. □ □

By Lemma 7 we see that for a tree T with n nodes and $u \in T$:

$$\overline{a}(u) - a(u) + 1 \leqslant |T_u| \, (1+\varepsilon)^{\lfloor \lg |T_u| \rfloor} \leqslant n \cdot 2^{\lfloor \lg n \rfloor \lg(1+\varepsilon)} \leqslant n \cdot 2^1 = 2n.$$

In particular, for any $u \in T$ we see that $a(u) \leqslant 2n$, and by Lemma 6 the function Assign-New is well-defined.

We are now ready to describe the labeling scheme:

Description of the encoder: Given an n-node tree T rooted in r, the encoding algorithm works by first invoking a call to Approx-New$(r, 0)$. Recall that by Lemma 6 we find $b(u)$ such that $b(u) - a(u) + 1 = \left\lfloor (1+\varepsilon)^k \right\rfloor$ as well as the value of k in $O(1)$ time. For a node u, denote the value of k by $k(u)$ and let x_u and y_u be the bit strings representing $a(u)$ and $k(u)$ respectively, consisting of exactly $\lceil \lg(2n) \rceil$ and $\left\lceil \lg(4 \lceil \lg n \rceil^2) \right\rceil$ bits (padding with zeroes if necessary). This is possible since $a(u) \in [2n]$ and $k(u) \in \left[4 \lceil \lg n \rceil^2 \right]$.

For each node $u \in T$ we assign the label $\ell(u) = x_u \circ y_u$. Since

$$\lceil \lg(2n) \rceil = 1 + \lceil \lg n \rceil, \quad \left\lceil \lg(4 \lceil \lg n \rceil^2) \right\rceil = 2 + \lceil 2 \lg(\lceil \lg n \rceil) \rceil = 2 + \lceil 2 \lg \lg n \rceil,$$

the label size of this scheme is $\lceil \lg n \rceil + \lceil 2 \lg \lg n \rceil + 3$.

Description of the decoder: Let $\ell(u)$ and $\ell(v)$ be the labels of the nodes u and v in a tree T. By the definition of the encoder the labels have the same size and it is $s = z + \lceil 2 \lg z \rceil + 3$ for some integer $z \geqslant 1$. By using that $s - \lceil 2 \lg s \rceil - 3 = z - O(1)$ we can compute z in $O(1)$ time. We know that the number of nodes n in T satisfies $\lceil \lg n \rceil = z$. We can therefore define ε to be the unique solution to $\lg(1+\varepsilon) = \lceil \lg n \rceil^{-1} = z^{-1}$. Let x_u and y_u be the first $z+1$ bits and last $\lceil 2 \lg z \rceil + 2$ bits of $\ell(u)$ respectively. We let a_u and k_u be the integers in $[2^{z+1}]$ and $[4z^2]$ corresponding to the bit strings x_u and y_u respectively. We define s_u as $\left\lfloor (1+\varepsilon)^{k_u} \right\rfloor$ and $b_u = s_u + a_u - 1$. We define a_v, b_v, k_v, s_v analogously. The decoder responds `True`, i.e. that u is the ancestor of v, iff $a_v \in [a_u, b_u]$.

Theorem 2 is now achieved by using the labeling scheme described above. Correctness follows from Theorem 1.

References

[1] Serge Abiteboul, Stephen Alstrup, Haim Kaplan, Tova Milo, and Theis Rauhe, *Compact labeling scheme for ancestor queries*, SIAM J. Comput. **35** (2006), no. 6, 1295–1309.

[2] Serge Abiteboul, Peter Buneman, and Dan Suciu, *Data on the web: From relations to semistructured data and XML*, Morgan Kaufmann, 1999.

[3] Serge Abiteboul, Haim Kaplan, and Tova Milo, *Compact labeling schemes for ancestor queries*, Proc. 25th ACM-SIAM symposium on Discrete algorithms, 2001, pp. 547–556.

[4] Stephen Alstrup, Philip Bille, and Theis Rauhe, *Labeling schemes for small distances in trees*, SIAM J. Discret. Math. **19** (2005), no. 2, 448–462.

[5] Stephen Alstrup, Cyril Gavoille, Haim Kaplan, and Theis Rauhe, *Nearest common ancestors: A survey and a new distributed algorithm*, ACM PRESS, 2002, pp. 258–264.

[6] Stephen Alstrup, Esben Bistrup Halvorsen, and Kasper Green Larsen, *Near-optimal labeling schemes for nearest common ancestors*, SODA, 2014, pp. 972–982.

[7] Stephen Alstrup, Haim Kaplan, Mikkel Thorup, and Uri Zwick, *Adjacency labeling schemes and induced-universal graphs*, CoRR **abs/1404.3391** (2014), To appear in STOC'15.

[8] Stephen Alstrup and Theis Rauhe, *Improved labeling scheme for ancestor queries*, Proc. 13th ACM-SIAM symposium on Discrete algorithms, 2002, pp. 947–953.

[9] ———, *Small induced-universal graphs and compact implicit graph representations*, FOCS '02, 2002, pp. 53–62.

[10] Edith Cohen, Haim Kaplan, and Tova Milo, *Labeling dynamic xml trees*, SIAM Journal on Computing **39** (2010), no. 5, 2048–2074.

[11] Alin Deutsch, Mary F. Fernández, Daniela Florescu, Alon Y. Levy, and Dan Suciu, *A query language for XML*, Computer Networks **31** (1999), no. 11-16, 1155–1169.

[12] Pierre Fraigniaud and Cyril Gavoille, *Routing in trees*, ICALP '01, Springer, 2001, pp. 757–772.

[13] Pierre Fraigniaud and Amos Korman, *Compact ancestry labeling schemes for xml trees*, In Proc. 21st ACM-SIAM Symp. on Discrete Algorithms (SODA, 2010.

[14] Pierre Fraigniaud and Amos Korman, *An optimal ancestry scheme and small universal posets*, STOC '10, STOC '10, 2010, pp. 611–620.

[15] Cyril Gavoillea, David Peleg, Stéphane Pérennesc, and Ran Razb, *Distance labeling in graphs*, Journal of Algorithms **53** (2004), 85–112.

[16] Sampath Kannan, Moni Naor, and Steven Rudich, *Implicit representation of graphs*, SIAM Journal On Discrete Mathematics, 1992, pp. 334–343.

[17] Michal Katz, Nir A Katz, Amos Korman, and David Peleg, *Labeling schemes for flow and connectivity*, SIAM Journal on Computing **34** (2004), no. 1, 23–40.

[18] Amos Korman, *Labeling schemes for vertex connectivity*, ACM Transactions on Algorithms **6** (2010), no. 2.

[19] Robert Endre Tarjan, *Finding dominators in directed graphs*, SIAM J. Comput. **3** (1974), no. 1, 62–89.

[20] Mikkel Thorup and Uri Zwick, *Compact routing schemes*, SPAA '01, 2001, pp. 1–10.